Reverse-Time Tracking to Enhance Passive Sonar

Garfield R. Mellema
Defence R & D Canada – Atlantic
P.O. Box 1012, 9 Grove Street
Dartmouth, NS, Canada B2Y 3Z7
garfield.mellema@drdc-rddc.gc.ca

Abstract - Passive sonar depends on signals of opportunity to detect, track and localize targets. These signals are typically detected and then tracked using Kalman filter-type signal followers. Target motion analysis (TMA) is then used to estimate the target's range and, from this, its position, course and speed. The accuracy of TMA is strongly dependent on the duration of the available track. Initiating a second tracker in reverse time at the time of detection can reduce or eliminate the delay between target detection and localization. A detection and tracking system for a passive sonar using a towed array receiver is described and an example of reverse-time tracking using real data is provided.

Reverse-time tracking is able to significantly increase the amount of track data that can be extracted from already available data, highlighting the need for improved data fusion. Potential improvements to this enhanced system through track association are discussed.

Keywords: Detection, tracking, localization, track association, track fusion, reverse-time tracking, retrodiction, sonar, passive sonar, towed array.

1 Introduction

Passive sonar uses acoustic signals of opportunity to detect, track and localize targets. Its subtle and covert nature allows an operator to efficiently and inconspicuously maintain situational awareness in the maritime environment [1].

There is a wealth of information encoded in the received signals; information about the sources of those signals, the paths along which the signals travelled, and the things that they met along the way, but that information requires significant patience and skill to retrieve. To answer the challenge, passive sonar processing is becoming more automated and, as the cost of computer processing and storage decreases, more services and options are becoming available, either as operator aids or as autonomous systems [2].

There are several stages of sonar data refinement, from the raw hydrophone data to target signature, position, course and speed, and the sonar picture can be examined at any of them. Not all data is suitable for refinement, however, and so at each stage only some of the data can be refined and advanced to the next level of representation. As a result of this cascading effect, only the smallest portion of the original data is represented at the highest level. Improvements in the data extraction and refinement processes also have a ripple effect, making more data available at each of the subsequent levels as well, making improvements beneficial at all stages of the refinement process.

The conversion of acoustic intensity maps into signal tracks occurs early in the passive sonar processing sequence and so it is an excellent candidate for improvement. A detection process searches for sequential features at a consistent bearing and frequency in the received data and, when the detection criteria are met, initiates a signal follower, typically a Kalman filter-type tracker, such as a Probabilistic Data Association Filter (PDAF), to refine the sequence of features into a track indicating bearing, frequency and relative signal power [3]. These tracks, under the right conditions, can subsequently be converted into estimates of target position, course and speed.

Traditionally, automated signal tracking has been done only in forward time, with all tracks updated simultaneously as each new set of data is acquired from the receiver. Initiating a second, asynchronous tracker in reverse time from those same detections has not been reported. This can be attributed to some extent to the cost and complexity of both randomly accessing prior time steps from the large quantity of recent historical data and processing track updates out of sequence.

Reverse-time tracking can improve the quantity of tracks available from the existing acoustic intensity data. Since both the forward and reverse-time tracks can be assumed to be tracking the same signal, an immediate increase in track duration can be seen when both track segments are associated across time. This increase in track duration benefits the subsequent target motion analysis (TMA) process [4]. The added track segments also provide additional opportunities for track association across frequency, which consequently improves opportunities for target identification. A third benefit, enhanced identification of track segments in the acoustic intensity plots, reduces the amount of ambiguous data remaining at this lower level of refinement.

2 The passive sonar data refinement process

In a typical towed array passive sonar system the received data is refined in a number of distinct stages as shown in

Report Documentation Page	*Form Approved OMB No. 0704-0188*

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
JUL 2006		**00-00-2006 to 00-00-2006**

4. TITLE AND SUBTITLE	
Reverse-Time Tracking to Enhance Passive Sonar	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Defence R & D Canada ? Atlantic,P.O. Box 1012, 9 Grove Street,Dartmouth, NS, Canada B2Y 3Z7,**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

9th International Conference on Information Fusion, 10-13 July 2006, Florence, Italy. Sponsored by the International Society of Information Fusion (ISIF), Aerospace & Electronic Systems Society (AES), IEEE, ONR, ONR Global, Selex - Sistemi Integrati, Finmeccanica, BAE Systems, TNO, AFOSR's European Office of Aerospace Research and Development, and the NATO Undersea Research Centre. U.S. Government or Federal Rights License

14. ABSTRACT

see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**8**	

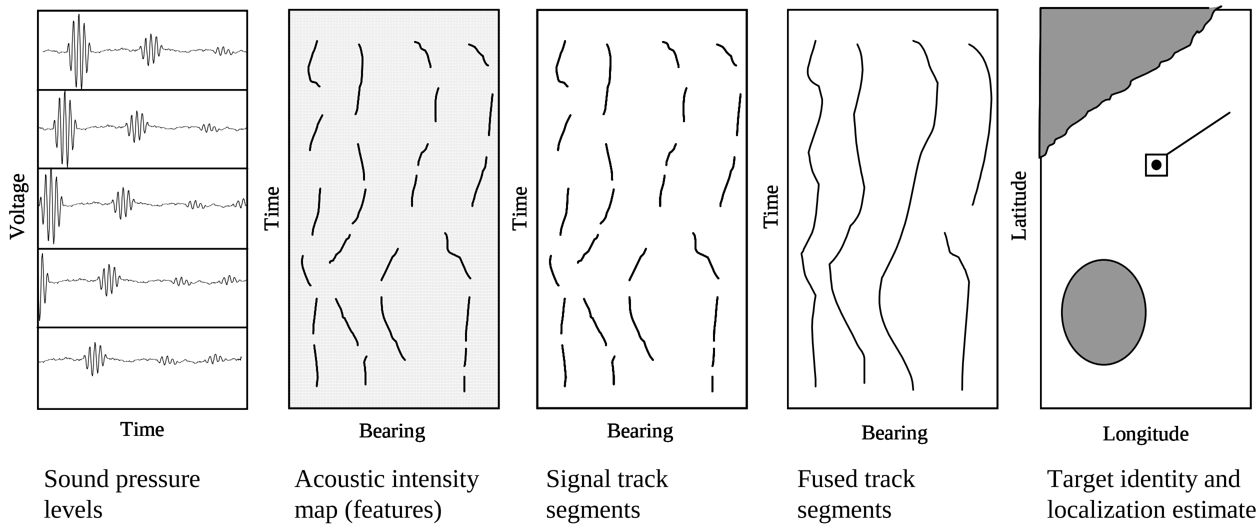

Figure 1. The sonar data refinement process can be divided into five stages.

Figure 1. As the data is advanced through these stages its information content is increased, although not all data at each level is suitable for advancement.

In the first stage, sound pressure levels in the water are converted into electrical signals and digitized by a linear series of hydrophones in the towed array receiver. The display here shows amplitude versus time at each omnidirectional hydrophone element.

By phase shifting and summing the contributions from each receiver, the array output can be steered to maximize sensitivity in a given direction. The result of this beamforming is a 2-dimensional map of received intensity versus bearing over time as shown in the second panel. This result is also available as a 3-dimensional map including frequency as an axis, although this is often displayed as a series of 2-dimensional plots, called beammaps, covering regularly spaced bearing intervals. Significant intensity peaks in these beammaps are described as features. Continuity of the features in bearing and/or frequency has not been established at this point, although it may be apparent to the eye.

In the third stage, detection and tracking algorithms are used to identify sequences of features in the beammaps and record these as tracks. Each track contains a history of the bearing, frequency and intensity of a signal. The continuity of each record is broken whenever contact with the signal is lost and so a single intermittent signal may produce a series of independent track segments.

It should be noted here that any signals that are not detected and tracked at this stage cannot be refined into track segments. They continue to be available for analysis as intensity peaks, but will not be eligible for any further processing or analysis at higher levels. Conversely, any spurious tracks such as those produced from noise or interference will be eligible for further processing or analysis at higher levels. It is therefore very important that the detection and tracking processes be both sensitive and selective.

At this level, each of these track segments contains only a single dimension of positional information, relative bearing, although it has a history of this information over time. TMA combines this track information with a best guess of the target's course and speed to estimate its range. Alternately, tracks provided by two separated observers can be used without guesses to estimate the position, course and speed of a target. In both cases, since the reliability of the TMA estimate is strongly dependent on the duration of the track being analyzed, it is strongly beneficial to precede TMA with track association.

A single observer could produce two sets of bearing information sequentially by changing course after acquiring the first set of bearing information. Using the forward-time tracking only approach, this may require about half an hour between initial detection and initial TMA estimate. This delay is due to signal propagation issues, low update rate, low SNR and the time required to reorient the sensor.

In the fourth stage, the tracks produced by the detection and tracking process are analyzed and grouped with respect to the source and/or platform from which they are believed to have originated. This is beneficial in several ways. First, the fusing of track segments across time into master tracks extends the duration of tracks available for TMA. Second, the association of tracks from the same source arriving at differing bearings, and therefore along distinct paths, can be used to estimate the range of the source by triangulation. Third, since sources and platforms can often be identified by the set of frequencies at which they radiate acoustic energy, the association of track segments across frequency provides a means for target identification.

In the final stage, TMA can be applied to refine those tracks having suitable characteristics into estimates of target position, course and speed. By assuming that the target is following a fixed course at constant speed, an estimate of the target's range can be made by analyzing the bearing rate of the target's track. Observations made before and after a change in the observer's course can be used to cross-fix the target. The accuracy of the estimates improves with the duration of the track. If multiple tracks at differing frequencies can be associated with the same target an estimate of identity can also be made. This highly refined format greatly facilitates the exchange of target information with other interested parties.

3 Typical passive sonar detection and tracking

3.1 Overview

The purpose of signal detection and tracking is to identify and recover target information from the non-specific information available from a sensor, such as a towed array sonar receiver. The process can be divided into 5 stages.

1. Pre-process the streams of acoustic pressure values from the multiple sensors in the towed array into a matrix of received acoustic intensity versus bearing and frequency over time and make them available through a data server, locally known as the beammap server.
2. Identify features in the matrix of received acoustic intensities that might indicate the presence of an incoming signal from an acoustic source.
3. Identify a sequence of features to increase confidence that the features represent the presence of an acoustic source. Assume that the source of the features is relatively stable and continuous.
4. Initiate a tracking algorithm to search for and incorporate additional features representing the signal, using them to build a track history representing the signal.
5. Terminate the track when there is no further reliable information about the signal.

3.2 Pre-processing

The acoustic receiver used in this system is a towed array, consisting of a number of omnidirectional hydrophone sensors assembled in a neutrally buoyant tube that can be towed through the water by a ship. The output of each sensor is digitized and recorded. The directional response of the receiver can be steered by phase shifting and summing the signals received at the sensors, a process known as beamforming.

The result of the beamforming calculation is a beam, or vector, of received acoustic intensity versus frequency at a given time and bearing. The noise floor at each frequency in each beam is determined independently and subtracted off, using either the mean or the median in a local frequency window centred on that point. When fully assembled across all bearings, the data set is a 3-dimensional matrix of acoustic intensities over frequency, bearing and time. Time-slices of the matrix are described as beammaps and are provided to the detection and tracking processes through a beammap server. Use of a beammap server permits both online and offline processing. Non-acoustic data, such as array heading, speed and depth are also available from this server.

3.3 Detection and track initiation

Beammaps are processed synchronously by the detection and tracking system, on a time-step-by-time-step basis. The sensitivity of the detection process is specified through the probability of false alarm, P_{fa}. This value, along with the noise distribution of the current beammap is used to determine the threshold above which intensity peaks will be considered as features.

The noise distribution of a beammap is a probability distribution of the acoustic power in its peaks, where a peak is defined as a cell having acoustic power greater than that in any of the eight adjacent cells. The detection level is the ($1-P_{fa}$) percentile level in the noise distribution. Only those peaks greater than the detection threshold are considered to be features. Those below the threshold are discarded as noise.

The results of peak detection in sequential beammaps are compared by an auto-detector to identify sequences of features at a similar bearing and frequency over time. Features that fall within the gate of an already active PDAF are ignored.

When a sequence of features is identified, the auto-detector initiates a simple, recursive alpha tracker to follow the sequence in order to determine whether it is sufficiently reliable to justify declaring a detection and initiating a new PDAF. The alpha tracker updates the position of a bearing-frequency gate into which subsequent features in the sequence are expected to fall based on the position of the features observed up to that time.

The criterion for declaring the detection of a signal is that *m* sequential features fall within the bearing-frequency gate within *n* time steps. The values of m and n are configurable and are typically based on the noise level of the overall set of data. When the *m* out of *n* criterion is met, a detection is declared and a PDAF is initiated.

The only information provided to the PDAF by the detector is the location of the feature closest to the frequency and bearing state values of the alpha tracker at the time of the detection. In the next time-step the newly established tracker will begin searching for its first track value with the assumption that the target has not moved significantly since its detection. It is left to the new tracker to determine the frequency and bearing rates.

Each feature can be evaluated no more than once. Features adjacent to existing PDAF tracks cannot be used by alpha trackers and features cannot be shared by alpha trackers. If an alpha tracker does not locate its m^{th} feature within the allotted *n* time steps, the alpha tracker is terminated and all of its data is discarded.

3.4 Tracking

Once the presence of an acoustic signal has been detected, a PDAF is initiated to continue tracking the signal. Continuous contact with these signals will allow the user to maintain awareness of the presence and quantity of local targets as well providing information that can be used to localize and identify them. In some cases, the only reliable contact with a target may be acoustic.

The measurement data is available as time-slices, or beammaps, from the beammap server. These are fed into the forward-time tracking process on a time-stepped basis and all of the trackers updated synchronously.

Since the tracking algorithm has some idea of what it is looking for and where to look for it, the size of the region to be processed can be limited by the predicted position of the signal being tracked and the maximum closing rate of the anticipated acoustic sources. Therefore, only the relevant window of beammap data is processed for each tracker. The windowing is handled within the forward-time tracking process.

The tracking process is intended to be more precise than the detection process and the peak detection process feeding the trackers is therefore also more precise. Three-point parabolic interpolation is used to improve the bearing and frequency resolution of the peak detection where possible, while two-point interpolation is used at the edges of the beammap window with the assumption that the peak lies between the two points. The interpolated peak values are described as Local Maximum Data (LMD). The same detection threshold is used in the LMD processing as in the auto-detection peak processing, although the peak detection processes are otherwise independent.

3.5 Track termination

Once initiated, the PDAF's will continue to run until a track termination condition is reached and a loss-of-lock is declared. The trackers evaluate five specific loss-of-lock criteria, each of which reflects an abnormal tracking state.

1. When the state bearing rate exceeds the maximum closing rate of the anticipated acoustic sources on a sustained basis.
2. When the state frequency rate exceeds the maximum expected Doppler rate on a sustained basis.
3. When no LMD values are found in the ordinary or expanded gate of the tracker on a sustained basis.
4. When a significant oversupply of LMD values is found in the gate of the tracker on a sustained basis.
5. When, on a sustained basis, the locations of the LMD values in the gate of the tracker are sufficiently distant from the predicted location that the β_0 value of the PDAF becomes overly large.

When any of these criteria for loss of lock is met, the tracker is terminated. The final tracker report includes a field specifying the reason for the termination.

4 Reverse-time tracking

4.1 Introduction

In our passive sonar tracking problem, the data available to us is a map of intensity versus bearing and frequency, continuously updated over time. The detection process searches for indications of repetitive activity over time, anywhere in the data. Once a repetitive event has been detected, the process changes to tracking the activity by predicting where it will next occur and then searching for evidence of the activity at that location.

The predictive nature of the tracking process is part of what makes it more sensitive to the observation of repetitive events than a simple detection process. It is therefore reasonable to expect that a target could have been tracked prior to the time at which it was detected, had the tracking process already been in effect at that time. The constraint we have imposed on the current detection and tracking process is that we expect the detection to precede the tracking although there is no necessary requirement for it to do so.

In many, if not most, applications of detection and tracking, the data storage and processing costs are significant and the acquired data is highly time-sensitive. This is not necessarily the case for bearings-only sonar data. Reliable TMA, especially in the single observer case, requires a minimum duration of bearing history, which, under the usual detect-then-track scheme, means an additional delay before a target can be localized. One method of reducing this delay would be to extend the tracks initiated by the detector as far back in time as possible.

The significant processes related to the production, propagation and scattering of acoustic signals are reversible in time. It is reasonable therefore to initiate and run signal followers in reverse as well as forward time. A signal follower running in reverse time could easily recover the short sequence of potential track in the detection window and, owing to the more robust character of a signal follower relative to the detection process, also follow the detected signal backwards through time prior to the start of the detection window. While there is an aesthetic appeal to maximizing the conversion of features into tracks, there is also the very practical benefit of reducing the target localization time and generally improving situational awareness.

4.2 Concept

Consider the scenario shown in Figure 2. An *m* out of *n* detector is evaluating features in the received data and eventually comes upon a sequence of features that meet the criteria for a detection. A detection is signalled at time t_0 and a PDAF is initiated. It follows the signal until some later time, t_1, when the conditions for loss-of-lock are met and the tracker is terminated.

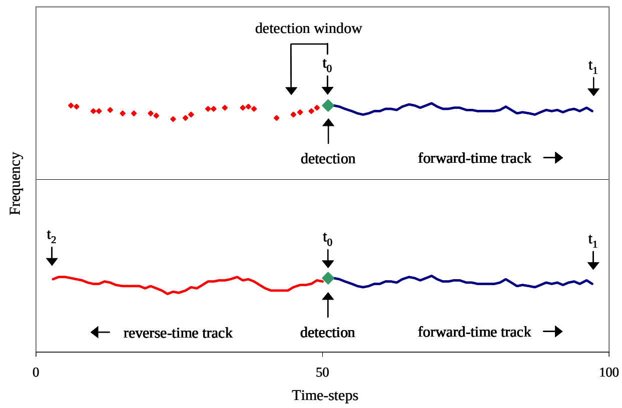

Figure 2. A typical detection and tracking scenario. A 5-out-of-7 detector initiates a forward-time tracker (upper) or both forward-time and reverse-time trackers (lower).

Although there is some evidence of a target prior to the detection, the existence of the target is not signalled until sufficient evidence has been accumulated. In hindsight, however, once the detection has been made, a review of the prior data shows clear evidence of the target, certainly enough to justify a detection decision. The usual tracker running in forward time uses none of this data. It is left as ambiguous data in the acoustic intensity plot.

Applying a second PDAF to data acquired prior to the detection can clarify the acoustic picture through increased conversion of features into target tracks. In the second panel a second signal follower is initiated when the detection is made, but this tracker is fed data in reverse-time order. The tracker follows the signal until some prior time, t_2, when the conditions for loss-of-lock are met and the track is terminated.

4.3 Implementation issues

The enhancement of a detection and tracking system to include reverse-time tracking raises a number of interesting issues.

Forward-time tracking is typically synchronous, such that all tracks are updated at the same time as each new time-step of data becomes available. Reverse-time tracking is necessarily asynchronous and requires significant interaction with the storage system as it requests beammaps in reverse-time order. The additional loading can be mitigated by storing computed values, such as the noise distribution and the detection threshold with each beammap and limiting each tracker request to only that portion of the beammap in which its next track update is likely to be found.

When a signal is detected and a forward-time tracker is initiated, the only information provided to the tracker is the location of a single feature in the acoustic intensity plot. It is left to the tracker to determine the bearing and frequency rates. The tracker determines these rates by examining features in a gate around the initiation point at subsequent time steps. This is true in both the forward-time and reverse-time cases. There is a small risk therefore that, if these subsequent features are sufficiently skewed, the initial forward-time and reverse-time tracks might not follow reciprocal paths. This is more likely to be an issue if the signal is just at the threshold of detectability.

Misalignment of the forward-time and reverse-time track segments is especially problematic if the pair of tracks originating from each detection are fused by default. It could be mitigated by a more complex initiation strategy. One possibility is to initiate one tracker and then, when it has stabilized, use its values to initiate the other. This, however, leaves open the possibility that the first tracker might terminate before stabilizing and the second tracker, however viable, might never be initiated. A better option is to transfer more information from the detection process to the trackers on initiation, ensuring that both trackers are initiated under similar conditions. It is also consistent with the idea that time is reversible.

A pair of trackers following the same signal in forward and reverse time can be expected to produce similar but not identical tracks. The differences are primarily due to smoothing delay and tracker response delay. The former is the result of applying a boxcar average to their state bearing and frequency rates. The latter is due to the fact that the trackers anticipate the future based on the past and the apparent past is different in forward and reverse time. A good example of this is a course change by the target.

The detection process is synchronized with and monitors the forward-time tracking process to avoid initiating new tracks where a tracker is already running. This is, however, not possible in the reverse-time case.

Consider the case of a signal that is repeatedly detected and then lost as it fades in and out. This signal might be represented by multiple forward-time and reverse-time tracks, where each forward-time track terminates prior to the initiation of the following forward-time track. Although the detection process monitors the list of currently active forward-time trackers to ensure that signals currently being tracked are not redetected, no such protection exists for the reverse-time trackers. Since the termination criteria might not be met simultaneously in both forward and reverse time, it is possible for a reverse-time tracker to overrun a forward-time track, continue following a reliable signal, and then overrun the corresponding reverse-time track. Additional reverse-time trackers, initiated at later times, might overrun multiple prior reverse-time tracks.

The existence of multiple track segments following the same signal is not necessarily disadvantageous and can be subsequently resolved through track association. The simple expedient of terminating a track as it begins to overrun another is a poor choice since it unnecessarily restricts the duration of the track segments, a highly valued property, and significantly undermines the opportunity to subsequently associate the two tracks.

5 An example of reverse-time tracking

The benefits of reverse-time tracking are best demonstrated using real data from the kind of system in which it would be applied. Figure 3 shows data from a sea trial on the Scotian Shelf using an SQR-19 towed array sonar receiver. The data was processed using the detection and tracking system described in Section 3 and then reprocessed using an enhanced version of the same system that initiated both forward-time and reverse-time trackers at each detection. The resulting tracks are shown in the upper and lower panels respectively.

The detection criteria used here is 9 out of 10 and the P_{fa} for both the detector and trackers is 0.001. These values were chosen to ensure that only a limited number of highly stable tracks were produced. The data began at $t = 32$ seconds and the update interval was 8 seconds. The detectors were started at $t = 40$ seconds.

The earliest detections were at $t = 192$ seconds and the first forward-time trackers were initiated at $t = 200$ seconds. A total of 15 detections were made and all of the forward-time trackers ran without loss of lock until the system was stopped at $t = 512$ seconds.

The lower frame shows the result of processing the same data using the enhanced detector-tracker configuration. In all but two cases the reverse-time trackers were able to run without loss of lock until they ran out of data at $t = 32$ seconds. The use of reverse-time

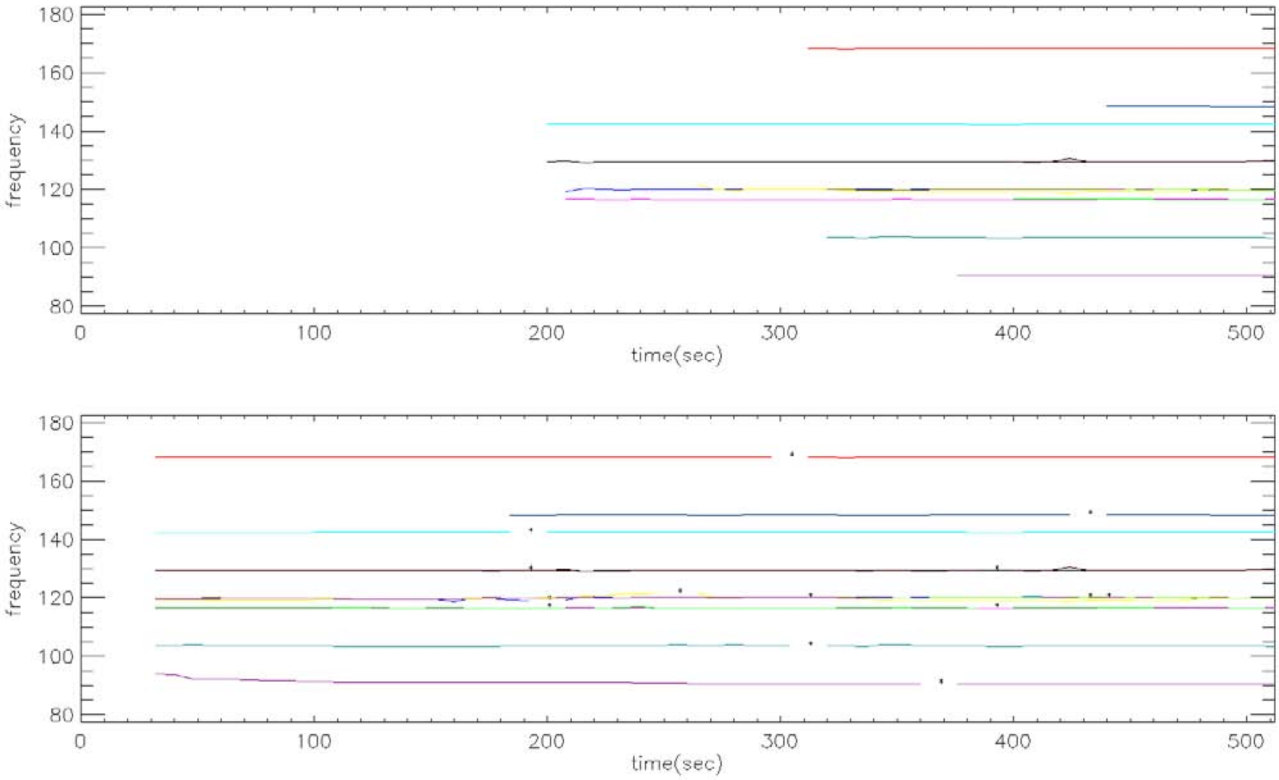

Figure 3. An example of forward and reverse-time tracking using field data. The upper panel shows only the forward-time tracks, the lower panel shows both. Asterisks mark the detections.

tracking has, in all cases, extended the available data by at least 104 seconds and as much as 400 seconds. The tracker at 148 Hz, which did lose lock, ran from t = 424 seconds until t = 184 seconds before losing lock due to due to an excessive number of LMD values in the gate. It extended the available data by 240 seconds.

Comparing the upper and lower frames it is clear that the signals being tracked existed well before they were detected and that this lag is due to more than just the minimum of 72 seconds required to trigger the m out of n detector. Clearly the detector has had to reset its count along the way due to missing features, although this does not appear to have been a problem to the reverse-time trackers. The additional track data they provide is due to more than just the detection delay.

6 Discussion

6.1 Increased track data

The example shown in the previous section clearly shows that reverse-time tracking can significantly increase the amount of track data that can be extracted from existing feature data. Initiating an additional tracker in reverse time whenever a traditional forward-time tracker is started immediately doubles the number of tracks produced. The increase in the total duration of all track segments is dependent on the characteristics of the feature data, but is clearly significant in this typical case.

In this example, most of the reverse-time trackers were terminated when they ran out of feature data at the beginning of the recording. While this demonstrates that these tracks easily continued well past the start of the detection window, it provides little information as to how far past. The answer to this is strongly dependent on the qualities of the signal being tracked. Since the forward-time and reverse-time trackers use the same tracking algorithm, it is reasonable to expect that, given a randomly chosen starting point on a continuous signal with a relatively constant signal-to-noise ratio (SNR), the tracks produced in each direction would be, on average, approximately equal in length. In the situation under consideration here, however, the starting point is not randomly chosen. Tracks start whenever a new detection is signalled. We therefore know only that the SNR of the underlying signal was barely sufficient to trigger the detector immediately prior to the time of detection.

It is possible however that in the more distant past the SNR of the underlying signal was sufficient to trigger the detector and temporarily sustain a forward-time tracker (as in the case of an intermittent signal). In this case the reverse-time tracker only needs to maintain lock on the newly detected signal until is reaches the termination point of the prior forward-time tracker. Once the reverse-time tracker reaches this point it should have little difficulty following the signal back to the start of the previous detection window and possibly beyond. Experience has shown that it is not unusual for a reverse-time tracker to overrun a prior forward-time track or, more precisely, for portions of a signal to be represented as both forward-time and reverse-time tracks. This situation is extremely beneficial as this pair of tracks can be easily and reliably associated and fused.

The example also contains two reverse-time tracks that terminated normally, due to loss of lock. This is not an unusual situation in the forward-time or reverse-time cases. They too provided additional track data.

6.2 Improved opportunities for TMA

The increased conversion of feature data into track data through the use of reverse-time tracking has a cascade effect on all subsequent stages of passive sonar signal development. It also has a very immediate effect on an operator's ability to do TMA.

TMA requires a minimum track duration in order to estimate the range, course and speed of a target. The accuracy of the estimate improves as the track duration increases. Reverse-time tracks can be developed without delay from existing feature data as soon as a detection is signalled. The duration of the immediately available reverse-time track may be sufficient to provide an immediate, initial estimate of target range, course and speed or possibly a more highly refined estimate. In any case, the length of forward-time track required to estimate target range, course and speed is reduced by the length of the reverse-time track, reducing the delay between detecting and localization a target.

The reverse-time track at 148 Hz that was terminated due to loss of lock, for example, could have reduced the delay between detection and localization by over 4 minutes. This is a very significant factor.

One of the advantages of active sonar is its ability to immediately localize targets. Its most significant drawback is its lack of stealth. The use of reverse-time tracking has the potential to combine the rapid localization capability of active sonar with the stealth of passive sonar.

6.3 An increased requirement for track fusion

Track association provides value to the passive sonar refinement process in three ways. First, it increases the duration of tracks by fusing related track segments into longer tracks. This improves continuity in the track picture and increases the opportunities for and accuracy of TMA. Second, track association reduces the number of independent tracks in the track picture, thereby reducing the perceived number of potentially independent targets. Third, associating tracks at multiple frequencies into a single signature provides an operator with a tool to identify target platforms.

With the implementation of reverse-time tracking making more track segments available, there are more opportunities and a greater need for track association of all kinds. The most pressing needs for association tools are in the areas of 1) association by a common detection, 2) association of concurrent tracks, and 3) association of tracks sharing a common origin.

If each detection initiates trackers in both forward and reverse time, it is essential that a mechanism be found to ensure that the two resulting tracks can be reliably fused. Originating from a common detection is insufficient by itself, as the two trackers are currently initiated with only a single instance of zero-order bearing and frequency information, i.e. a single interpolated peak. The trackers must then independently estimate first-order information such as bearing and frequency rates from features in the subsequent forward and reverse time steps respectively. Even in the most likely case, where both sets of features are due to the same signal, there is no guarantee that the track formed by the union of the two tracks will not change abruptly as it passes through the detection point.

This problem can be addressed by using higher order information to initiate the trackers. A potentially better method, though more complex, is to initiate one tracker from the other after it has stabilized. To avoid having to wait additional time steps for new data to be acquired, the forward-time tracker could be initiated from the reverse-time tracker, or the forward-time tracker could be initiated at the first feature used by the detector and its state vector used to initiate the reverse-time tracker at the end of detection window. Since the features in the detection window are already known, a third possibility would be to initiate the two trackers at either end of the detection window and feed them the features used by the detector as their first m data points. The type of mechanism to be used and its implementation is an area for further investigation.

In the traditional detection and tracking system, it is possible for multiple trackers to be following the same signal. The case can easily be made for the fusion of these tracks. These situations are unusual in forward time because the detection and tracking system was designed to provide an exclusion area around each tracker to avoid repeated detections of the same signal. In most of these cases, the existence of multiple concurrent trackers can be attributed to the presence of severe noise or interference.

The enhanced detection and tracking system described here provides no protection against a reverse-time tracker overrunning a prior forward-time tracker on the same signal or against detecting a signal that will later be followed by a reverse-time tracker. Implementing the latter case would be extremely complex, if not unworkable, and neither case is necessarily problematic. A reverse-time tracker overrunning a previously terminated forward-time track on the same signal is not only increasing the amount of available track data, it is also building evidence for the fusion of the two tracks. Once sufficient evidence has been established for that fusion the concurrent portion of the reverse-time tracker is redundant, although it will likely continue overrunning the forward-time track as well as the subsequent reverse-time track.

This is an excellent opportunity for track fusion as these track segments no longer represent unique signals by default and their fusion would reduce redundancy and produce a single track of greater duration than either of the constituent tracks. Reliable track fusion is, of course, crucial in this situation.

A target platform typically contains multiple acoustic sources and can often be identified as to type or hull number by the specific combination of acoustic signals that it emits. Passive sonar can be used to identify targets by associating multiple signals, represented as tracks, that are concurrent in time and bearing but differ in frequency. Reverse-time tracking, by increasing the rate at which signals are converted into tracks, increases the opportunities for target identification.

One of the simplest and most reliable types of track association is the association of tracks that differ only in frequency. These are tracks that have followed the same propagation path from the target to the receiver and will therefore appear at a similar bearing. They also share similar, if not proportional, changes in frequency over time. Sets of associated frequencies can be used to differentiate targets or compared with templates for identification.

When this type of association is employed, the length of time that a target can be continuously tracked increases from the duration of the longest track segment to the end-to-end duration of all of the fused track segments. As with the two previous cases, this increase in the length of continuous bearing track improves continuity in the track picture, as well as the opportunities for, and the accuracy of, TMA.

6.4 Storage and processing implications

In the prototype reverse-time tracking system, the increase in storage and computational load was significant relative to the original forward-time tracking only system, even with a relatively low P_{fa}. This is partially due to the ad hoc way in which the system was assembled, as an attachment to the original forward-time tracking only system. Regardless, the system ran at better than real time on a dual-processor 2.0 GHz Pentium 4 desktop computer.

A significant part of the loading was the recall from disk and transfer of beammaps from previous time steps. Much of the processing cost is due to the asynchronous nature of the reverse-time trackers, which means that each tracker independently requires a new beammap at each time step. Caching data in memory as well as on disk could significantly reduce the beammap retrieval delay.

There are often on the order of hundreds of synchronized forward-time trackers running at any time. Since reverse-time trackers are asynchronous and run to completion immediately on initiation, each represents a short burst of intense loading. In order to ensure that the pace of the forward-time trackers is not compromised, the reverse-time trackers should be run at a lower priority or executed on a separate computing platform. None of these storage or processing issues presents a significant impediment to the implementation of reverse-time tracking.

7 Conclusions

7.1 Summary

An enhanced detector-tracker system has been constructed and used to process passive towed array sonar data. A comparative example of the tracks produced by the traditional and enhanced systems is provided. The enhanced system works well and significantly improves the ability to extract information from the available sonar data.

The increased number of track segments improves opportunities for track association and target localization at subsequent levels of refinement and reduces the time required for target localization and identification. TMA requires a minimum track duration before a reliable localization estimate can be made. Since it is constructed from acoustic intensity data cached prior to the detection, a reverse-time tracker can immediately run to completion without delay once it has been initiated. This can provide several minutes or more of reliable track data that can be fused with the track produced by the traditional forward-time tracker and used for target localization and identification.

The implementation of reverse-time tracking requires minor changes to the traditional detector-tracker system, and the bulk of the additional processing could be hosted on an independent computing platform requiring only cached information from the traditional system.

The use of reverse-time tracking has the potential to combine the rapid localization capability of active sonar with the stealth of passive sonar.

7.2 Future work

Along with the increased conversion of features into tracks comes an increased requirement for track association, especially in the following areas.

In the enhanced detection and tracking system demonstrated here, the forward-time and reverse-time tracks initiated by each detection are treated as independent track segments. Their common initiation conditions however may be sufficient evidence to justify their association. This is probably best handled by improving the method by which the pair of forward-time and reverse-time trackers is initiated when a signal is detected.

An intermittent signal can be expected to produce some simultaneous forward-time and reverse-time tracks. Since these concurrent tracks represent the same signal, they should be fused.

Track association can also be used to identify concurrent track segments that originated from the same target and share common bearing data, but differ in frequency. These could be used to extend the duration of a target track for TMA or for target identification.

References

[1] R.J. Urick, Principles of Underwater Sound, 3rd ed., McGraw-Hill, New York, 1983.
[2] R.O. Nielsen, Sonar Signal Processing, Artech House, Boston, 1981.
[3] Y. Bar-Shalom and T.E. Fortmann, Tracking and Data Association, Academic Press, Boston, 1988.
[4] S. Rudnicki, *Target Motion Analysis using Frequency and Bearing Measurements*, Maritime Engineering Journal, July 1993.